

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 0 969 375 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:  
05.01.2000 Bulletin 2000/01

(51) Int Cl.7 G06F 11/32

(21) Application number: 99305159.8

(22) Date of filing: 30.06.1999

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventor: Stoutamire, David P.  
San Juan Bautista, California 95045 (US)

(74) Representative: Browne, Robin Forsythe, Dr.  
Urquhart-Dykes & Lord  
Tower House  
Merrion Way  
Leeds LS2 8PA West Yorkshire (GB)

(30) Priority: 30.06.1998 US 108049

(71) Applicant: SUN MICROSYSTEMS, INC.  
Palo Alto, California 94303 (US)

(54) Method for visualizing locality within an address space

(57) A method and that display certain aspects of a computer memory on a display screen in accordance with a fractal curve. In at least one embodiment, the fractal curve is a Hilbert curve. The fractal curve display can

be used as an aid in visualizing various types of information about a computer memory. Certain embodiments have a display that remains static, while other embodiments have a display that is dynamic and changes in real-time to reflect changes in memory.




FIG. 2

EP 0 969 375 A2

## Description

### BACKGROUND

[0001] This application relates to computers and computer systems and further relates to a method and apparatus for visually displaying various aspects of memory.

[0002] As computer systems become more complex, it is becoming more difficult for human beings to easily understand and grasp all the intricacies involved in any given computer system.

[0003] During execution, a computer system may move units of memory, such as pages or cache lines between units of hardware. As more complex caching systems are designed and as caching systems begin to have more and more levels, it is becoming more and more difficult for human beings to understand whether a given caching system is operating efficiently for the software executing on the computer system.

[0004] Software currently exists to aid a human being in visualizing certain aspects of a computer's memory, such as the performance impact of caching. These conventional software programs, however, display memory using a line-by-line display on a display device. For example, in such a line-by-line display, a pixel representing information about a first unit of memory is displayed on a first line in a first column of a display screen. A pixel representing information about a second unit of memory is displayed on a first line in a second column of the display screen. A pixel representing information about a last unit of memory may be displayed on a last line in a last column of a display screen, and so on. Pixels representing information about intermediate units of memory are displayed, one after the other, on consecutive rows and in consecutive columns within each row (or vice versa in consecutive columns and rows).

[0005] The term "locality of memory" refers to the location of data (e.g., variables) within memory in time and space. Accessing a variable closely in time is called "temporal locality." Accessing variables stored physically close together is called "spatial locality." Physical locations in a memory that are close together have a high spatial locality. For example, a variable that is read from close together in time has a high temporal locality. For example, a 49<sup>th</sup> memory unit and a 50<sup>th</sup> memory unit (such as bits, bytes, words, pages, etc.) that are physically close to each other in a computer memory have a high spatial locality. As another example, if a variable is read twice in quick succession, the variable has a high temporal locality. On modern computer systems, temporal and spatial locality impact performance.

[0006] Unfortunately, in conventional line-by-line visualization software, even though a 49<sup>th</sup> memory unit is physically contiguous to a 50<sup>th</sup> memory unit in the actual memory of the computer, it is entirely possible that the display pixels corresponding to the memory units will be on opposite sides of the display screen. For example,

the pixel corresponding to the 49<sup>th</sup> memory unit might be the last column in row n, while the pixel corresponding to the 50<sup>th</sup> memory unit might be displayed in the first column of row n+1. Because conventional visualization tools do not necessarily show memory units having a high locality as being physically near each other on the display, conventional visualization tools are not ideal for aiding a human being in visualizing how various memory units relate to each other or for visualizing locality.

[0007] Another example of conventional visualization software displays memory accesses by address on one axis of a graph and time on the other axis. This method also fails to present data in a way that effectively aids human understanding of the relationships between the data.

### SUMMARY OF THE INVENTION

[0008] The described embodiments of the present invention provide a method and apparatus that display certain aspects of a computer memory on a display screen in accordance with a fractal curve instead of a line-by-line display. In at least one embodiment, the fractal curve is a Hilbert curve. The fractal curve display can be used as an aid in visualizing various types of information about a computer memory. Certain embodiments have a display that remains static, while other embodiments have a display that is dynamic and changes in real-time to reflect changes in memory.

[0009] For example, the display can be used as an aid to visualize memory access patterns, such as when and/or where read or write operations occur in a computer memory. As a second example, the display can be used as an aid to visualize which pages are currently being swapped in or out of various levels of memory in a virtual memory system. As a third example, the display can be used as an aid to visualize which cache lines are active in a caching system. As a fourth example, the display can be used as an aid to visualize which pages are being swapped into memory. As a fifth example, the display can be used as an aid to visualize the read/write permissions of various pages of memory. As a sixth example, the display can be used as an aid to visualize the shared/modified exclusive states of cache lines.

[0010] Use of a Hilbert curve as an aid in visualizing memory has the advantage that it allows memory groups equal to powers of two to be displayed near each other. For example, display bits of the fractal curve corresponding to memory addresses 00xxxxxx, where xx is any value, are grouped in a same quadrant of the memory display. Similarly, display bits of the fractal curve corresponding to memory address 0000xxxx are grouped within a sub-quadrant of the quadrant holding display bits for the addresses 00xxxxxx. Display bits of the fractal curve corresponding to memory address 000000xx are grouped within a sub-sub-quadrant of the quadrant holding display bits for the addresses

0000xxxx, and so on.

[0011] Certain embodiments use color on the display to indicate types of memory access, while other use display color to indicate a type of memory being represented in the display. In addition, certain embodiments of the present invention also use a 3-D version of a fractal curve, either drawn in 2D or rendered in 3D to aid in visualization of various aspects of memory. In accordance with the purpose of the invention, as embodied and broadly described herein, the invention relates to a method of displaying information, performed by a data processing system, comprising the steps of: receiving notification of a memory event, including a memory location; determining a portion of a fractal curve that corresponds to the memory location at which the memory event has occurred; and highlighting a portion of a display screen in accordance with the determined portion of the fractal curve.

[0012] A fuller understanding of the invention will become apparent and appreciated by referring to the following description and claims taken in conjunction with the accompanying drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0013] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate several embodiments of the invention and, together with the description, serve to explain the principles of the invention.

[0014] Fig. 1 is a block diagram of a computer system in accordance with one embodiment of the present invention.

[0015] Fig. 2 shows an illustration of a curve used by a computer program of Fig. 1 as an aid to visualizing certain aspects of a computer memory of Fig. 1.

[0016] Fig. 3 is a block diagram showing an overview of steps performed by an embodiment of the invention to map portions of memory onto a fractal curve.

[0017] Fig. 4 is a flow chart showing steps performed by an embodiment of the present invention to map portions of memory onto a fractal curve.

[0018] Fig. 5 shows a state table used by the flow chart of Fig. 4 to determine which portion of the fractal curve corresponds to a particular portion of memory.

[0019] Fig. 6 shows possible rotation values for the state table of Fig. 5.

[0020] Fig. 7 shows an example of levels of quadrants used in the flow chart of Fig. 4.

[0021] Figs. 8 and 9 provide an example showing how to locate a portion of a displayed fractal curve corresponding to a memory address.

[0022] Figs. 10(a) and 10(b) provide examples of three-dimensional fractal curves shown on a two-dimensional display.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

### I. General Discussion

[0023] Fig. 1 is a block diagram of a data processing system 100 in accordance with a preferred embodiment of the present invention. In Fig. 1, data processing system 100 includes a processor 102 and a data storage area (e.g., a memory) 104. The example of Fig. 1 shows that processor 102 contains an exemplary cache 111, although not all systems on which the present invention is implemented contain a cache and not all caches in such systems are necessarily located in processor 102. Data storage area 104 includes certain well-known types of data, such as data structures, pages, page tables, etc. 106. Storage area 104 also includes memory visualization software 110, which aids in visualization of various aspects of the computer system, as described below in detail. Memory visualization software 110 could also be located in a computer system different from system 100, such as a remote system (not shown). Storage area 104 preferably also includes software (not shown) for communicating with a network, such as a LAN, WAN, or the internet, although the invention also can be implemented on standalone computers.

[0024] System 100 preferably includes an input device 112, such as a keyboard, pointing device, mouse, touchpad, etc., which allows input to be read by processor 102. System 100 preferably also includes an output device 140. This output device can be, for example, a computer monitor, a display device, or any device that allows processor 102 to send or display output.

[0025] A person of ordinary skill in the art will understand that system 100 may also contain additional elements, such as input/output lines, input devices, such as a keyboard, a mouse, and a voice input device; and display devices, such as a display terminal. System 100 may also include a computer readable input device 150, such as a floppy disk drive, CD ROM reader, or DVD reader, that reads computer instructions stored on a computer readable medium 160, such as a floppy disk, a CD ROM, or a DVD disk. System 100 also may include application programs, operating systems, data, etc., which are not shown in the figure for the sake of clarity. It also will be understood that system 100 can also include numerous elements not shown, such as disk drives, keyboards, display devices, network connections, additional memory, additional processors or CPUs, LANs, input/output lines, etc.

[0026] In the following discussion, it will be understood that the steps of methods and flow charts herein discussed herein preferably are performed by processor 102 (or another appropriate processor) executing instructions stored in storage area 104 (or other appropriate memories or storage areas). It will also be understood that the invention is not limited to any particular implementation or programming technique and that the

invention may be implemented using any appropriate techniques for implementing the functionality described herein. The invention is not limited to any particular programming language or operating system.

[0027] The instructions in storage area 104 may be read into storage area 104 from the computer-readable medium 160. Execution of sequences of instructions contained in main memory causes one of the processors to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0028] The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to a processor for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as a storage device. Volatile media includes dynamic memory. Transmission media include coaxial cables, copper wire and fiber optics, including the wires that comprise a bus within a computer. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0029] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, a hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertapes, any other physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave, or any other medium from which a computer can read.

[0030] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to a processor for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to the computer system can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector coupled to a bus can receive the data carried in the infra-red signal and place the data on the bus. The bus carries data to main memory, from which a processor retrieves and executes the instructions. The instructions received by main memory may optionally be stored on a storage device either before or after execution by a processor.

## II Visualization of Memory

[0031] Fig. 2 shows a curve 200 used by a computer program of Fig. 1 as an aid to visualizing certain aspects of a computer memory of Fig. 1. In the described em-

bodiment, the curve is not drawn explicitly. Instead, memory addresses of, for example, storage area 104 or cache 111, are mapped onto the curve. Curve 200 can, in fact, correspond to any appropriate memory, whether in the computer 100 or a secondary storage 120. The display includes a fractal curve (such as a Hilbert curve), where each edge of the curve represents an associated portion or aspect of the memory. As explained below in more detail, use of a fractal curve to represent portions of or aspects of a memory allows a human being to more easily visualize relationships between the portions of memory.

[0032] Specifically, the use of a Hilbert curve allows each quadrant and subquadrant of the curve to correspond to respective bits of a memory address. The nature of a Hilbert curve allows easy visualization of the locality of memory/hardware, since addresses that reside on even powers of two map into quadrants and subquadrants of the fractal curve. Rectangular regions of the curve of a given size naturally correspond to hardware entities, such as cache lines or pages, so it is easy for a human being to see how memory access patterns map onto hardware. Portions of the display that are visually nested correspond to units of hardware that are logically nested. For example, in Fig. 2, curve 200 is used to aid in visualizing a ten bit address. Fig. 2 shows a portion 250 of curve 200 that corresponds to locations identified by a page (which has 256 locations). As a further example, Fig. 2 shows a smaller portion 252 of curve 200, nested within portion 250, that corresponds to a cache line (which has 32 locations).

[0033] Fig. 3 is a block diagram showing an overview of steps performed by an embodiment of the invention to map portions of memory onto a fractal curve. In the described embodiment, the fractal curve is not drawn explicitly.

[0034] The steps of Fig. 3 can be performed by memory visualization software 110 being executed by processor 102. As shown, in step 301, software 110 receives notification of a memory event at a certain address. A memory event can be, for example, a memory access such as a read or write or a cache access. A memory event can also be the simple receipt of a next one of sequential memory addresses if the software 110 is to map all addresses in a memory onto the curve. In the described embodiment, notification of a memory event is received from the operating system of the data processing system, although notification of a memory event can be received from any appropriate source. For example, the Solaris operating system, available from Sun Microsystems, Inc., allows executing programs to request that the operating system notify them of certain memory events. (Solaris is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries.) As another example, other operating systems may allow the software 110 to poll the operating systems for the occurrence of memory events.

[0035] Step 302 determines a position on the fractal curve in accordance with the received memory address bits. Details of step 302 are shown in more detail in Fig. 4. Step 304 highlights the portion of the fractal curve corresponding to the received memory address. This portion can be a single pixel or a group of pixels, depending on the resolution of the particular implementation of the invention. Thus, for example, if software 110 is designed to highlight a portion of the curve whenever a read access occurs, every read access of a new address will highlight another pixel on the curve.

[0036] Fig. 4 is a flow chart showing steps performed by an embodiment of the present invention to generate a fractal curve display, where the display indicates the localities of memory accesses. The flow chart shows details of step 302 and 304 of Fig. 3. Step 401 displays a fractal curve on the display. As discussed above, certain embodiments may not perform this initial step.

[0037] In step 402, software 110 receives notification of the occurrence of a memory event at a memory address and, in step 404, chooses a color to represent the event. In some embodiments, all memory events are indicated by a single color. In other embodiments, various types of memory events are indicated by different respective colors. For example, read accesses may be displayed in a first color, while write accesses may be displayed in a second color. As another example, all memory locations having a read permission may be displayed in a first color, while all memory locations having a read/write permission may be displayed in a second color. As a third example, all memory locations accessed a certain number of times per time unit may be displayed in a first color, while all memory locations accessed a different number of times per time unit may be displayed in a second color. As a fourth example, all memory locations filled from a cache may be displayed in a first color, while all memory locations filled from secondary storage may be displayed in a second color. As a fifth example, all memory locations in a cache that have been written back to secondary storage may be displayed in a first color, while all memory locations that have not been written back may be displayed in a second color.

[0038] Other variations using more than two colors will be apparent. As a sixth example, all memory locations having a shared, modified, or exclusive state are displayed in respective colors.

[0039] Steps 406 through 414 form a loop that determines which portion of the fractal curve corresponds to the received memory address. This loop preferably is repeated for each two bits of the received address. In step 406, an initial "rotation" is set to a value of "1" as described below in connection with Fig. 6. An initial rotation of "1" is chosen because, in the example, the Hilbert curve is initially oriented in this rotation. Other initial orientations are possible.

[0040] Step 403 gets the next two bits of the received address. In step 410, software 110 uses the rotation and

the two bits of the memory address to index into a state table 114 shown in Figs. 1 and 5 to yield a "next rotation" value and new x,y values. These x,y values are saved for each iteration of the loop until an x,y value has been saved for each two bits of the received address. At this point, in step 416, software 110 highlights a portion of the fractal curve, in the chosen color, at a position on the curve represented by the saved sequence of x and y values. The saved x values can be concatenated to form the x position of the position on the curve. The saved y values can be concatenated to form the y position of the position on the curve. Further details of step 416 are shown by the examples below.

[0041] Fig. 5 shows an example of the state table 114 used by the flow chart of Fig. 4 to determine which portion of a fractal curve corresponds to a particular portion of memory. The state table 114 in the described embodiment includes input values: a current rotation value and two address bits; and output values: two x,y bits and a next rotation value. As discussed above, each two bits of a memory address along with a current rotation value are used to determine output x,y values and a next rotation value. The sequence of x,y values are saved and used to highlight a portion of the fractal curve corresponding to the memory address.

[0042] Fig. 6 shows possible rotation values for the state table 114 of Fig. 5. A rotation value refers to the order of two bit values in a quadrant used to draw the fractal curve. For example, the "L" rotation value refers to an order 602 in which the two bit values 00, 01, 10, and 11 form a counterclockwise path, starting at a lower left quadrant. The "R" rotation value refers to an order 604 in which the two bit values 00, 01, 10, and 11 form a counterclockwise path, starting at an upper right quadrant. The other rotation values "U", "D", "I", "r", "u", and "d" are similarly shown.

[0043] Fig. 7 shows an example of levels 702, 704, 706, and 708 of quadrants used in the flow chart of Fig. 4. At a first level 702, the space in which the fractal curve is to be drawn is divided into four quadrants. Since, in the Figure, the rotation is "1", the quadrants are numbered clockwise from the top left as 00, 01, 10, and 11. As further shown, each quadrant in level 1 is broken into four level 2 sub-quadrants, each also numbered 00, 01, 10, and 11, the order of which depends on the rotation value associated with each sub-quadrant. As further shown, each sub-quadrant in level 2 is broken into four level 3 sub-sub-quadrants, each also numbered 00, 01, 10, and 11, the order of which depends on the rotation value associated with each sub-sub-quadrant. As further shown, each sub-sub-quadrant in level 3 is broken into four level 4 sub-sub-sub-quadrants, each also numbered 00, 01, 10, and 11, the order of which depends on the rotation value associated with each sub-sub-sub-quadrant.

[0044] Thus, in the example of Fig. 7, each level 4 sub-sub-sub-quadrant corresponds to one eight-bit memory address. Given an initial rotation and initial two

address bits, a state table similar to the state table of Fig. 5 can be used to determine which portion of the fractal curve (i.e., which sub-sub-sub quadrant) corresponds to a given memory address. Note that the memory address can represent a single bit, a byte, a word, a cache line, a page, or any other appropriate memory unit.

[0045] It will be understood that the example of Fig. 7 is provided by way of example and that other fractal curves can have other organizations and that different lengths of memory addresses will have different numbers of levels of subquadrants. For example, a sixteen bit address where each edge of a level of the curve represents two address bits will have eight levels.

[0046] Figs. 8 and 9 provide an example showing how to locate a portion of a displayed fractal curve corresponding to a memory address. In the example, software 110 determines a portion of a fractal curve corresponding to a memory address 10001110. Because the address has eight bits, four x,y pairs are determined and saved in one embodiment, a level N storage 130 in accordance with the state table of Fig. 5. Specifically, as shown in Fig. 8, an initial rotation of "l" and initial two bit address values of 10 result in next rotation values of l, U, r, and r and further result in x,y pairs of 11, 00, 10, and 00.

[0047] As shown in the example of Fig. 9, the saved x,y pairs identify a portion 902 of the fractal curve. Specifically in the example, the first rotation value "l" and the first two address bits 10 identify a bottom right quadrant in the first level. The second rotation value "l" and the second two address bits 00 identify a top left quadrant in the second level. The third rotation value "U" and the third two address bits 11 identify a top right quadrant in the third level. The fourth rotation value "r" and the fourth x,y pair 10 identify a top left quadrant in the fourth level. The top left quadrant is the curve portion 902 at position (10 8) which can be a single pixel or multiple pixels, and which corresponds to the received memory address.

[0048] Figs. 10(a) shows an example of a first level of a three-dimensional fractal curve. Fig. 10(b) shows an example of a two connected three-dimensional fractal curves. This curve can be displayed to any desired depth of levels and in either two dimensions or three dimensions (given a proper display device), as is known by persons of ordinary skill in the art. As will be appreciated by persons of ordinary skill in the art, a portion of a three-dimensional curve corresponding to a memory address can be determined using a state table similar to that of Fig. 5.

[0049] In summary, the present invention displays memory locations in accordance with a fractal curve, such as a Hilbert curve. Use of a fractal curve to map memory locations causes locations having a similar locality to be displayed near each other, thus improving the ability of a human being to visually appreciate the locality of memory events.

[0050] While the invention has been described in conjunction with a specific embodiment, it is evident that many alternatives, modifications and variations will be apparent to those skilled in the art in light of the foregoing description. For example, the entire display can be updated at once, instead of pixel by pixel. Accordingly, it is intended to embrace all such alternatives, modifications and variations as fall within the spirit and scope of the appended claims and equivalents.

## Claims

1. A method of displaying information, performed by a data processing system, comprising the steps of:

receiving notification of a memory event, including a memory location;

determining a portion of a fractal curve that corresponds to the memory location at which the memory event has occurred;

highlighting a portion of a display screen in accordance with the determined portion of the fractal curve; and

displaying the fractal curve on the display screen.

3. The method of claim 1, wherein the fractal curve is a Hilbert fractal curve.

4. The method of claim 1, wherein the memory location is a memory location that has either a READ permission or a READ/WRITE permission.

5. The method of claim 1, wherein the memory event either reads from the memory location or writes to the memory location.

6. The method of claim 1, wherein the memory location has been accessed at least at a predetermined frequency.

7. The method of claim 1, wherein the memory location is a location in a cache memory that has been filled from another memory or that has been written back to another memory, and wherein the memory location has one of a shared, modified, and exclusive state.

8. The method of claim 1, wherein the fractal curve is a three-dimensional fractal curve displayed in at least two dimensions, wherein a plurality of portions of the display screen corresponding to the fractal curve are highlighted at once.

9. The method of claim 1, wherein the display is dynamic and changes as corresponding memory events occur.

10. The method of claim 1 wherein the determining 5  
step includes the step of determining a portion of a  
fractal curve in accordance with a state table

10

15

20

25

30

35

40

45

50

55




FIG. 1






FIG. 2




FIG. 3

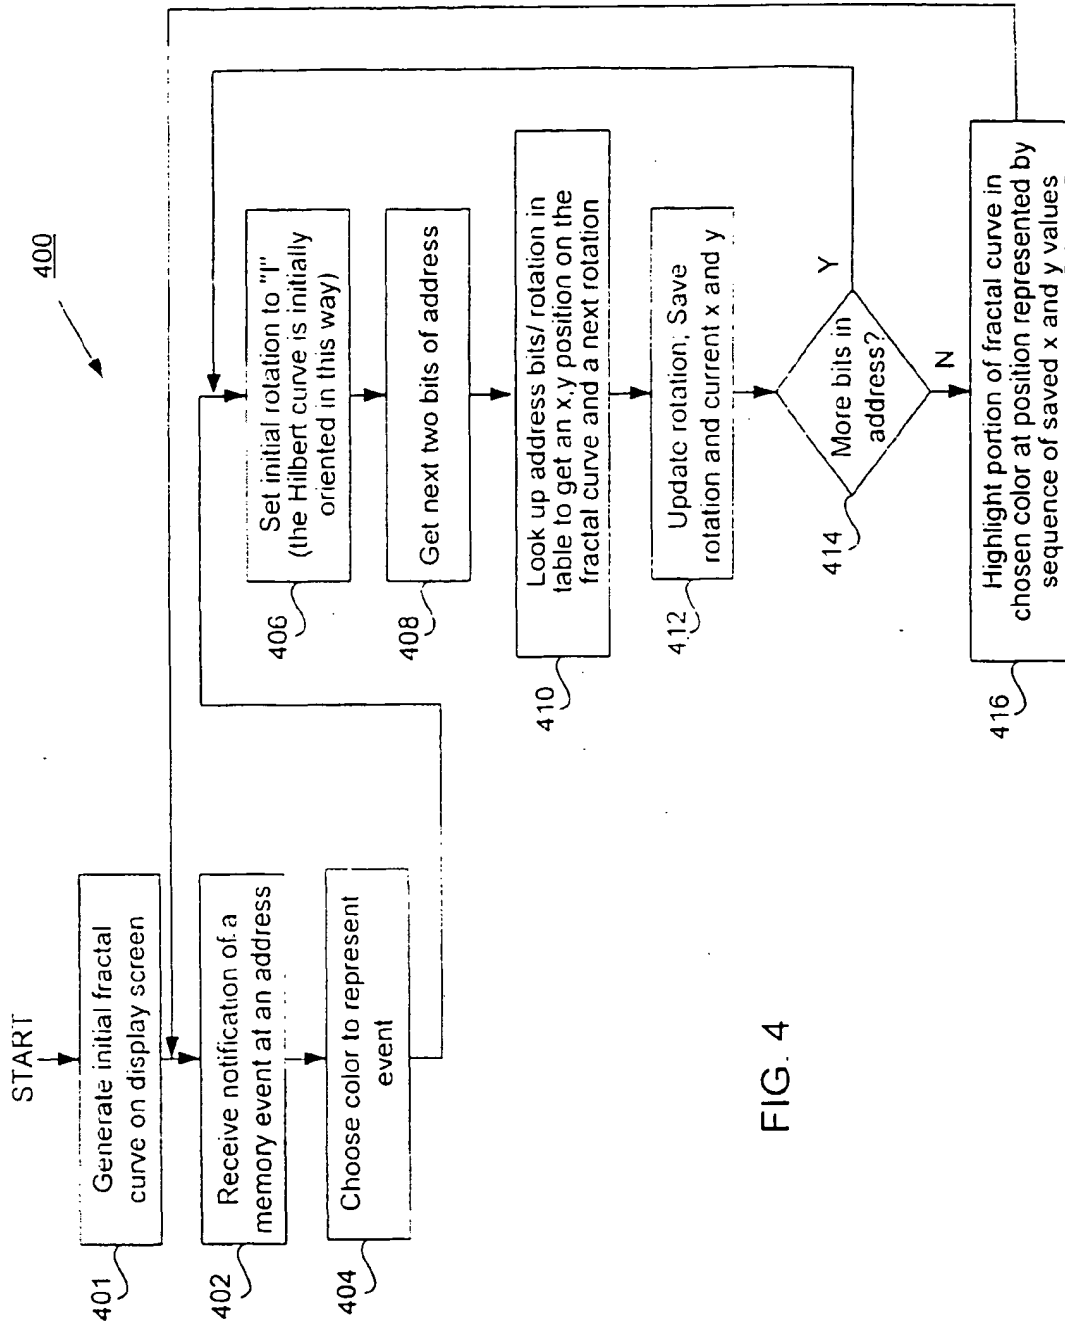


FIG. 4

Input (Current rotation)	Input (Address bits)	Output (x,y)	Output (next rotation)
I	00	0,0	U
I	01	1,0	I
I	10	1,1	I
I	11	0,1	D
L	00	0,1	d
L	01	1,1	L
L	10	1,0	L
L	11	0,0	u
r	00	1,1	D
r	01	0,1	r
r	10	0,0	r
r	11	1,0	U
R	00	1,0	u
R	01	0,0	R
R	10	0,1	R
R	11	1,1	d
u	00	1,0	R
u	01	1,1	u
u	10	0,1	u
u	11	0,0	L
U	00	0,0	I
U	01	0,1	U
U	10	1,1	U
U	11	1,0	r
d	00	0,1	L
d	01	0,0	d
d	10	1,0	d
d	11	1,1	R
D	00	1,1	r
D	01	1,0	D
D	10	0,0	D
D	11	0,1	I

FIG. 5




FIG. 6




FIG. 7

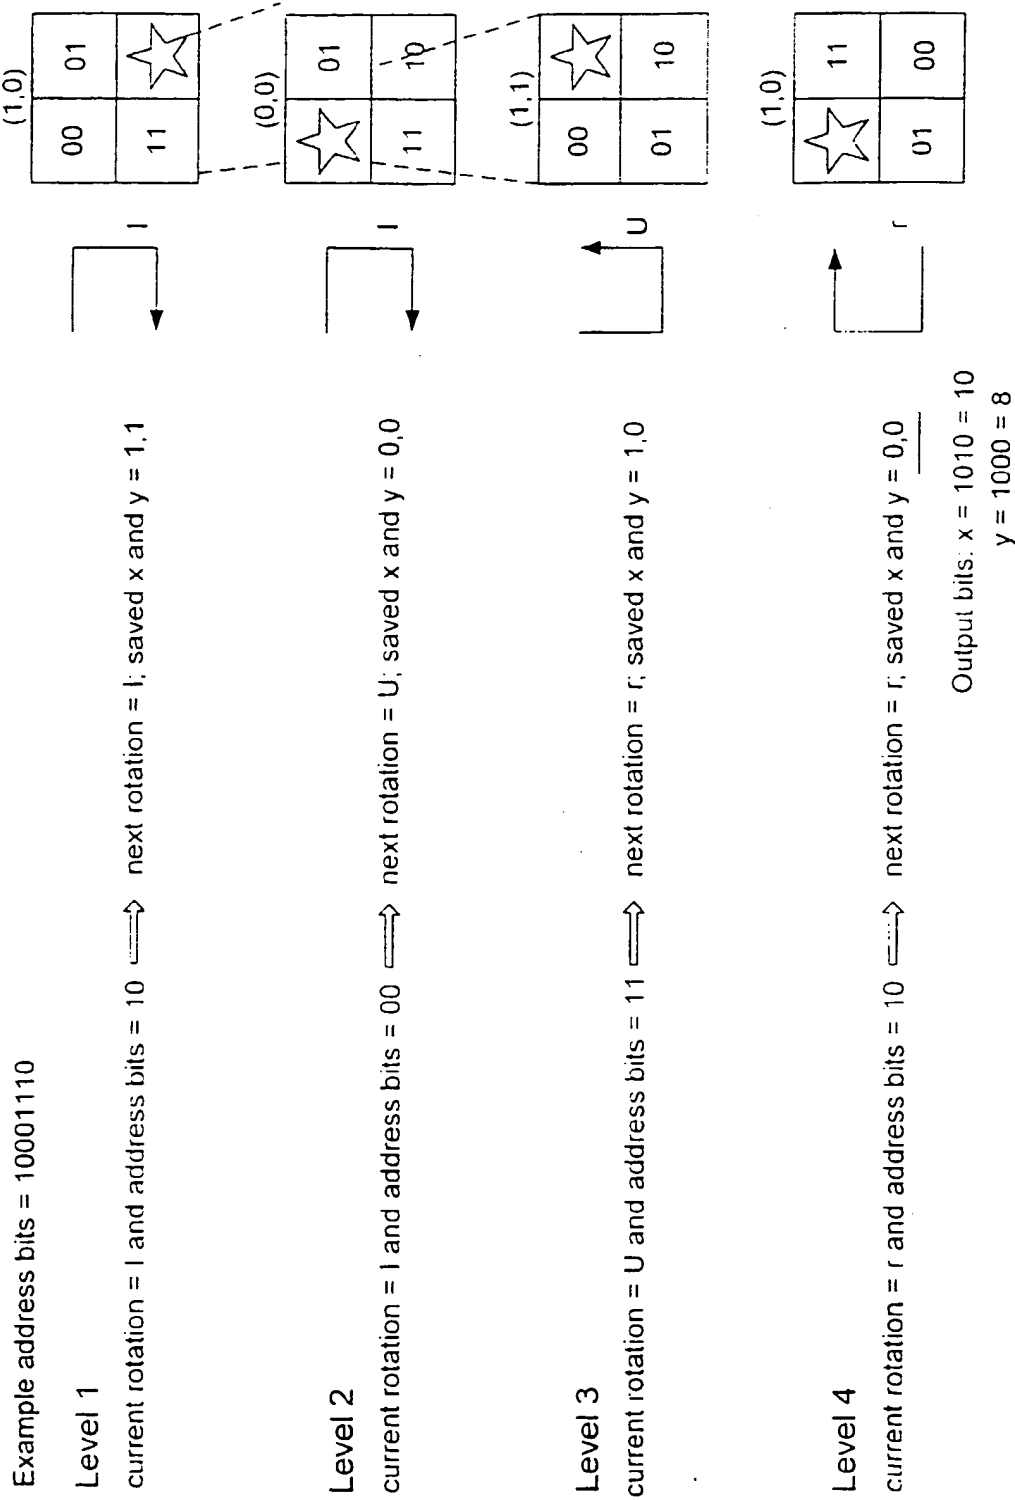


FIG. 8




FIG. 9




FIG. 10A




FIG. 10B